

SECURE REPOSITORY WITH LAYERS OF TAMPER RESISTANCE
AND SYSTEM AND METHOD FOR PROVIDING SAME

FIELD OF THE INVENTION

The present invention relates generally to the field of computer security and, more particularly, to a secure repository with layers of tamper resistance, and a system and method for providing such a repository.

BACKGROUND OF THE INVENTION

In the field of computer security, one enduring problem is to create a system that allows an owner of information to electronically distribute the information throughout the world while regulating use of that information on remote hardware over which the information owner has no control. For example, information may be delivered to an end user in encrypted form with the ultimate goal that it be viewed (but not copied or otherwise misused) by the end user. The information requires a key in order to be decrypted, but it may not be desirable to give the end user unfettered access to the key because the user could then copy the decrypted information and disseminate it at will.

One solution is not to provide the key directly, but to provide the key to the end user in the form of software that applies the key (or that performs some other sensitive function to be hidden from the user). Such software may contain various types of protection designed to complicate or resist attempts to analyze or misuse the software or the secrets that the software is designed to protect. However, the drawback to this solution is that attempts to create "secure" software incorporating such resistance have so far proven ineffective, as such software has invariably been misused, ported to unauthorized installations, or broken in one way or another. A further drawback is that if technology advances to permit greater protection to be built into the software, it is

1 not always possible to "renew" the security technology by replacing an old unit of
2 "secure" software with a new one.

3 In view of the foregoing, there is a need for a system that overcomes the
4 limitations and drawbacks of the prior art.

5 6 SUMMARY OF THE INVENTION

7 The present invention provides advantageous systems and methods for
8 creating and using a software-based secure repository. A black box provided herein is
9 an exemplary secure repository that uses cryptographic techniques, preferably
10 public/private key techniques, to perform decryption and authentication services in a
11 secure manner that resists discovery of secret keys used by the cryptographic
12 techniques.

13 A secure repository according to the invention, such as the exemplary
14 "black box" provided herein, functions as the trusted custodian of one or more
15 cryptographic keys. It performs cryptographic functions such as using a cryptographic
16 key to decrypt information. A user who wishes to use encrypted information is
17 provided with a black box that incorporates the key needed to decrypt the information.
18 The black box contains code that applies the key without actually representing the key
19 in memory, thus shielding the key from discovery by the user. Preferably, the
20 information is encrypted with a public key pair that is unique (or substantially unique)
21 to the user, and each user obtains a unique black box that applies that particular user's
22 private key. The black box may provide the decrypted information to other software
23 modules which the black box trusts not to misuse the information or to divulge it in an
24 unauthorized way. The black box may use cryptographic authentication techniques to
25 establish trust with these other software modules.

26 In order to obtain the black box, the user's computer contacts a black
27 box server, preferably via a network, and uploads a hardware identifier associated with
28 the computer. The black box server creates an "individualized" black box which

1 contains code to apply a particular cryptographic key, where the code (as well as other
2 code contained in the black box) is based on, and preferably bound to, the hardware
3 identifier. The black box server may introduce various types of protections into the
4 code, such as diversionary code, integrity checks, inline encryption, obfuscated
5 execution, code reorganization, and self-healing code. The black box server then
6 downloads an executable file or executable library containing the black box for
7 installation on the user's computer.

8 In a preferred embodiment, the black box interfaces with an application
9 program for which it provides secure functions by way of a decoupling interface that
10 makes the details of the black box transparent to the developer of the application
11 program. The decoupling interface may, for example, be an application programmer
12 interface (API) usable with multiple dynamic-link libraries (DLLs), where a different
13 DLL is linked to the application program at runtime depending on which black box is
14 being used. A new DLL may be created for a new black box that did not exist at the
15 time the application program was created. The use of a decoupling interface in this
16 manner supports a "renewable" model of security - i.e., the black box can be replaced
17 if it has been found to be defective or unsecure, or if later technological developments
18 permit the creation of an even more secure black box. The DLL that implements the
19 decoupling interface is an example of a software module that may be authenticated by
20 the black box.

21 Other features of the invention are described below.
22

23 **BRIEF DESCRIPTION OF THE DRAWINGS**

24 The foregoing summary, as well as the following detailed description of
25 preferred embodiments, is better understood when read in conjunction with the
26 appended drawings. For the purpose of illustrating the invention, there is shown in the
27 drawings exemplary constructions of the invention; however, the invention is not
28 limited to the specific methods and instrumentalities disclosed. In the drawings:

FIG. 1 is a block diagram showing an exemplary computing environment in which aspects of the invention may be implemented;

FIG. 2 is a block diagram showing an exemplary use of a preferred type of secure repository;

FIG. 3 is a block diagram showing a relationship between a computer that requests a secure repository and a computer that generates a secure repository;

FIG. 4 is a block diagram of an exemplary secure repository generator according to aspects of the invention;

FIG. 5 is a block diagram of a cryptographic code generator according to aspects of the invention;

FIG. 6 is a flow diagram showing an exemplary process for creating a machine-individualized secure repository;

FIG. 7 is a flow diagram showing an exemplary process for performing the code creation step of FIG. 6;

FIG. 8 is a block diagram showing an exemplary architecture that includes a decoupling interface for use with a secure repository and an application program;

FIG. 9 is a flow diagram showing an exemplary process of using a secure repository.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Overview

Modern computing and network technology has permitted widespread and low-cost electronic distribution of information. One feature of electronic information is that it is easily copyable and, once it has been delivered to a computer, the operator of that computer may use and disseminate the information in ways that are neither contemplated by, nor consistent with the commercial interests of, the owner of

1 the information. A secure repository may be used to control the use of information on
2 hardware over which the information's owner otherwise has no control.

3 Computing Environment

4 As shown in FIG. 1, an exemplary system for implementing the
5 invention includes a general purpose computing device in the form of a conventional
6 personal computer or network server 20 or the like, including a processing unit 21, a
7 system memory 22, and a system bus 23 that couples various system components
8 including the system memory 22 to the processing unit 21. The system bus 23 may be
9 any of several types of bus structures including a memory bus or memory controller, a
10 peripheral bus, and a local bus using any of a variety of bus architectures. The system
11 memory includes read-only memory (ROM) 24 and random access memory (RAM) 25.
12 A basic input/output system 26 (BIOS), containing the basic routines that help to
13 transfer information between elements within the personal computer 20, such as during
14 start-up, is stored in ROM 24. The personal computer or network server 20 may
15 further include a hard disk drive 27 for reading from and writing to a hard disk, not
16 shown, a magnetic disk drive 28 for reading from or writing to a removable magnetic
17 disk 29, and an optical disk drive 30 for reading from or writing to a removable optical
18 disk 31 such as a CD-ROM or other optical media. The hard disk drive 27, magnetic
19 disk drive 28, and optical disk drive 30 are connected to the system bus 23 by a hard
20 disk drive interface 32, a magnetic disk drive interface 33, and an optical drive
21 interface 34, respectively. The drives and their associated computer-readable media
22 provide non-volatile storage of computer readable instructions, data structures, program
23 modules and other data for the personal computer or network server 20. Although the
24 exemplary environment described herein employs a hard disk, a removable magnetic
25 disk 29 and a removable optical disk 31, it should be appreciated by those skilled in the
26 art that other types of computer readable media which can store data that is accessible
27 by a computer, such as magnetic cassettes, flash memory cards, digital video disks,

1 Bernoulli cartridges, random access memories (RAMs), read-only memories (ROMs)
2 and the like may also be used in the exemplary operating environment.

3 A number of program modules may be stored on the hard disk, magnetic
4 disk 29, optical disk 31, ROM 24 or RAM 25, including an operating system 35 (e.g.,
5 WINDOWS® 2000, WINDOWS NT®, or WINDOWS® 95/98), one or more application
6 programs 36, other program modules 37 and program data 38. A user may enter
7 commands and information into the personal computer 20 through input devices such as
8 a keyboard 40 and pointing device 42. Other input devices (not shown) may include a
9 microphone, joystick, game pad, satellite disk, scanner or the like. These and other
10 input devices are often connected to the processing unit 21 through a serial port
11 interface 46 that is coupled to the system bus 23, but may be connected by other
12 interfaces, such as a parallel port, game port, universal serial bus (USB), or a 1394
13 high-speed serial port. A monitor 47 or other type of display device is also connected to
14 the system bus 23 via an interface, such as a video adapter 48. In addition to the
15 monitor 47, personal computers typically include other peripheral output devices (not
16 shown), such as speakers and printers.

17 The personal computer or network server 20 may operate in a networked
18 environment using logical connections to one or more remote computers, such as a
19 remote computer 49. The remote computer 49 may be another personal computer,
20 another network server, a router, a network PC, a peer device or other common
21 network node, and typically includes many or all of the elements described above
22 relative to the personal computer 20, although only a memory storage device 50 has
23 been illustrated in FIG. 1. The logical connections depicted in Fig. 2 include a local
24 area network (LAN) 51 and a wide area network (WAN) 52. Such networking
25 environments are commonplace in offices, enterprise-wide computer networks,
26 Intranets and the Internet.

27 When used in a LAN networking environment, the personal computer or
28 network server 20 is connected to the local network 51 through a network interface or

1 adapter 53. When used in a WAN networking environment, the personal computer or
2 network server 20 typically includes a modem 54 or other means for establishing
3 communications over the wide area network 52, such as the Internet. The modem 54,
4 which may be internal or external, is connected to the system bus 23 via the serial port
5 interface 46. In a networked environment, program modules depicted relative to the
6 personal computer or network server 20, or portions thereof, may be stored in the
7 remote memory storage device 50. It will be appreciated that the network connections
8 shown are exemplary and other means of establishing a communications link between
9 the computers may be used.

10 Exemplary Use of a Secure Repository

11 Referring now to FIG. 2, there is shown an exemplary use of a secure
12 repository in connection with remote computer 49. Remote computer 49 is a computer
13 that may be able to receive information from other computers, for example by means of
14 wide-area network 52 (which may be or comprise the network known as the Internet),
15 local-area network 51, or physical delivery on media such as removable magnetic disk
16 29 or optical disk 31. Remote computer 49 typically is not in the control or dominion of
17 the owner/provider of such information, and thus a secure repository, such as black box
18 240, may be used to control the use of information on a computer 49 over which the
19 owner/provider of the information otherwise has no control.

20 Black box 240 is an exemplary embodiment of a secure repository. Black
21 box 240 comprises code which executes on the general processing means provided by
22 remote computer 49 (which may be analogous to the processing means of computer 20
23 depicted in FIG. 1, including, but not limited to, processing unit 21). Black box 240 is
24 preferably equipped with one or more cryptographic keys 248, and contains code to use
25 the cryptographic keys to perform on or more cryptographic functions, such as
26 decryption of encrypted information 204, and/or authentication of cryptographically
27 signed data 212. In an exemplary embodiment, cryptographic keys 248 include
28 asymmetric key pairs for use with public/private-key cryptographic methods.

1 Preferably, black box 240 also contains code that is designed to ensure that black box
2 240 performs its cryptographic functions only for trusted software, and in an
3 environment that resists divulgence of sensitive results or attacks on black box 240
4 itself. It will be appreciated by those of skill in the art that decryption and
5 authentication are merely exemplary, and not limiting, of the type of secure or sensitive
6 functions that could be performed by a secure repository, such as black box 240.
7 Secure repositories that perform other functions may be installed or employed on
8 remote computer 49, without departing from the spirit and scope of the invention.

9 Application program 244 is a software program that also executes on
10 remote computer 49. FIG. 2 depicts an example where black box 240 performs
11 decryption and/or authentication services for application program 244, which is a
12 program that in some way processes or uses information. In this example, application
13 program 244 provides encrypted information 204 to black box 240. Black box 240, in
14 turn, decrypts encrypted information 204 (e.g., by using one or more cryptographic
15 key(s) 248) and returns decrypted information 208 to application program 244.
16 Similarly, application program 244 may call upon black box 240 to authenticate an item
17 of data 212, and black box 240 may determine the authenticity of data 212 (e.g., by
18 using one or more cryptographic key(s) 248) and may optionally provide to application
19 program 244 an indication 216 of whether data 212 is authentic. Encrypted information
20 204 may, for example, be textual information (e.g., a novel), digital audio (e.g. music),
21 digital video (e.g., a movie), financial data, software, or any other type of information
22 (confidential or otherwise). Data 212 to be authenticated may comprise a signed
23 certificate or other signed information. In a typical use, data 212 includes a certificate
24 attesting to black box 240 that application program 244 is a sufficiently trustworthy
25 program (i.e., that application program 244 can be trusted to handle decrypted
26 information 208 without violating rules applying to the usage of decrypted information
27 208). Data 212 could also comprise a signed "license" to use decrypted information
28 208, and black box 240 (or trusted decoupling interface 220 (see below)) may

1 authenticate the signed license to verify that the license is not actually a forgery that has
2 been proffered by a user to gain unauthorized access to decrypted information 208.
3 Depending on the nature of encrypted information 204, application program 244 may
4 be a text-rendering program (i.e., a program that enables the reading of electronically-
5 distributed encrypted text by displaying the text as print characters on a monitor), a
6 video-rendering program, an audio-rendering program, or any other type of program
7 that is capable of processing decrypted information 208 in some manner (or that, for
8 any reason, needs to authenticate cryptographically signed information). It will be
9 appreciated that the arrangement shown in FIG. 2 is particularly advantageous, because
10 it allows application program 244 to make use of cryptographic services (i.e.,
11 decryption and authentication), without having to know the details of how those
12 cryptographic services are performed, or the particular key(s) 248 that are used to
13 perform them.

14 Associated with remote computer 49 are one or more hardware IDs 224.
15 A hardware ID 224 comprises data that identifies, or otherwise relates to, particular
16 hardware associated with remote computer 49. Preferably, hardware ID 224 comprises
17 one or more of the following component IDs: CPUID 228, processor serial number
18 232, and hard disk ID 236. CPUID 228 may be a unique (or substantially unique)
19 pseudo-random number assigned to a central processing unit (CPU) (not shown) of
20 remote computer 49 by the manufacture of such CPU. CPUID 228 may be retrievable
21 by executing an instruction on such CPU that provides the CPUID to a running
22 program. Processor serial number 232 may be a sequentially-assigned number
23 associated with a CPU (or other processor) of remote computer 49, which may also be
24 retrievable by executing an instruction on such CPU. Some processors may have both a
25 CPUID 228 and a serial number 232; other processors may have only one but not the
26 other. Hard disk ID 236 may be a number assigned (e.g., pseudo-randomly or
27 sequentially) to a hard disk (not shown) associated with remote computer 49. Hard disk
28 ID 236 may be assigned by the manufacturer or distributor of such hard disk, and may

1 be stored in an accessible location on such hard disk. A function of hardware ID 224 is
2 to provide a number that is uniquely (or substantially uniquely) associated with remote
3 computer 49. Preferably, hardware ID 224 is based on all of the following component
4 IDs: CPUID 228, serial number 232, and hard disk ID 236. However, hardware ID
5 224 may be based on a subset of those component IDs, or on entirely different
6 component IDs (which may related to different hardware of remote computer 49, or,
7 perhaps, to the serial number of an operating system (not shown) installed on remote
8 computer 49). Moreover, hardware ID 224 may comprise the entirety of the component
9 IDs on which it is based (e.g., a concatenation of the component IDs), a subset of those
10 component IDs (e.g., the rightmost sixteen bits from each component ID), or may be
11 derived by a function that generates hardware ID 224 based on the component IDs
12 (e.g., the multiplicative product of a plurality of component IDs modulo 2^{32}). It is also
13 possible to use one of the component IDs itself as a hardware ID 224; for example,
14 hardware ID 224 could simply be equal to CPUID 228. Preferably, hardware ID 224
15 cannot easily be retrieved or derived in an environment other than remote computer 49,
16 such that a program can be constructed that relies on retrieving hardware ID 224 from
17 the execution environment during execution, thus binding the program to remote
18 computer 49. Hardware ID(s) 224 may be created/determined in any manner that in
19 some way relates to the hardware and/or environment of remote computer 49, without
20 departing from the spirit and scope of the invention. Additionally, remote computer 49
21 may have a plurality of hardware IDs 224, with each hardware ID 224 being based on
22 different component IDs, or on a different function of the same component IDs.

23 Black box 240 may access hardware ID(s) 224, and may use hardware
24 ID(s) 224 in order to perform computations relating to its functions. For example, in
25 the course of performing a function such as decryption, black box 240 may need to use
26 a number n . Instead of storing n directly, black box 240 may store n minus the value of
27 a hardware ID 224. In this case, in order to use the number n , black box 240 would
28 contain code that that retrieves or computes hardware ID 224 and adds it to the stored

1 value to create n . This has the effect of “binding” black box 240 to a particular remote
2 computer 49, since it needs a value that can only be accessed on remote computer 49.
3 As a simple example, if hardware ID 224 is equal to CPUID 228, black box 240 could
4 execute an instruction to retrieve CPUID 228 and add it to the stored value. In this
5 case, the instruction will not produce the correct value n unless black box 240 is
6 executing on the remote computer 49 that has the correct CPUID 228.

7 It should be noted that black box 240 is designed not merely to perform
8 sensitive functions in a hidden and protected manner, but also to ensure that the
9 sensitive results that it produces cannot escape to an untrusted software module – which
10 means in practice that black box 240 performs the function of authenticating other
11 software modules that are part of the secure environment in which black box 240
12 operates. Preferably, black box 240 authenticates every such software module. In
13 general, each such software module maintains, in a known place, a signed hash of its
14 code, where the signature is generated with a private key. Cryptographic keys 248 may
15 include the corresponding public key which may be used by black box 240 to verify
16 such signatures and establish trust with those software modules.

17 Black box 240 may communicate directly with application program 244,
18 or, alternatively, it may communicate through decoupling interface 220. Decoupling
19 interface 220 is essentially an intermediary by way of which black box 240 and
20 application program 244 may communicate using a common language or protocol. The
21 use of decoupling interface 220 may improve the versatility of a distribution system that
22 uses black box 240, since it allows application program 244 to communicate with
23 different types of black boxes 240, and allows black box 240 to communicate with
24 different types of application programs 244. Additionally, decoupling interface 220 may
25 authenticate application program 244 to black box 240. As discussed in further detail
26 below, one issue that arises in the use of a black box 240 that is separate from the
27 application program 244 for which it provides services is that black box 240 should not
28 perform sensitive functions for application program 244 unless it is certain that

1 application program 244 can be trusted to use the sensitive functions only in approved
2 ways, which means in practice that application program 244 must satisfy the protocol
3 and other requirements necessary to proffer a trust certificate to black box 240. When
4 decoupling interface 220 is used, application program 244 need only be able to
5 communicate with decoupling interface 220, where various embodiments of decoupling
6 interface can be provided each of which can authenticate application program 244 (and
7 decoupling interface 220 itself) to a particular black box 240. Decoupling interface 220
8 may, for example, be provided by the supplier of black box 240 to facilitate
9 communication between black box 240 and application program 244, as well as to
10 facilitate the “renewability” of black box 240 (as more particularly discussed below). It
11 will be appreciated by those skill in the art, however, that while decoupling interface
12 220 is a convenient and advantageous structure, communications between black box
13 240 and application program 244 may take place either directly or through decoupling
14 interface 220 without departing from aspects of the invention.

15 Provision/Acquisition of Black Box 240

16 Remote computer 49 acquires black box 240 from a black box server
17 304. Black box server 304 may be implemented on a typical computing device, such as
18 computer 20 (shown in FIG. 1). (Hereinafter, black box server 304 and the computer
19 20 on which it is implemented shall be referred to interchangeably, unless context
20 indicates otherwise.) Referring now to FIG. 3, black box server 304 is preferably
21 connected to remote computer 49 via wide-area network 52. Preferably, wide-area
22 network 52 is or comprises the network known as the Internet. Black box server 304
23 includes a black box generator 37a. Black box generator 37a is preferably a computer
24 program running on computer 20 (such as one of the “other programs” 37 depicted in
25 FIG. 1). Black box generator 37a receives a hardware ID 224 of a remote computer 49,
26 and generates a black box 240 that is “individualized” for remote computer 49.
27 “Individualized” in this context means that the contents of a given black box 240 is at
28 least partly based on the hardware ID associated with the remote computer 49 for which

1 the black box is created, and that a first black box 240 created for a first remote
2 computer 49 is different from (or, at least, is very likely to be different from) a second
3 black box 240 for a second remote computer 49. The various black boxes 240 are
4 "different" in the sense that they contain different cryptographic keys 248, different
5 code to apply the keys 248, and different obfuscating code. The particular ways in
6 which the black boxes 240 can be made different are discussed more particularly below
7 in connection with FIG. 4.

8 In a typical use of black box server 304, remote computer 49 contacts
9 black box server 304 and transmits to black box server 304 a request for a black box
10 240 via wide-area network 52. The request for a black box may be generated as part of
11 a registration or "activation" process relating to a particular type of software. For
12 example, application program 244 may be delivered without a black box 240 that is
13 needed for its use or some aspect thereof. For example, application program 244 may
14 be able to work with unencrypted information (or even "keyless sealed" information) in
15 the absence of black box 240, but may need black box 240 in order to be "activated"
16 for use with encrypted information 204. In this case, application program 244 may
17 include activation software 308 which is invoked upon the first use of application
18 program 244, where activation software 308 contacts black box server 304 in order to
19 obtain a black box 240 which will enable use of application program 244. As another
20 example, software 308 may be general-purpose web-browsing software by means of
21 which a user of remote computer 49 may issue a request to black box server 304 by
22 navigating to a site operated by black box server 304.

23 The request for a black box 240 that comes from remote computer 49
24 preferably includes the hardware ID(s) 224 of remote computer 49. Upon receiving the
25 request, black box server 304 invokes black box generator 37a and proceeds to create a
26 black box 240 that is at least partly based on hardware ID(s) 224, and is thus
27 "individualized" for remote computer 49. As discussed above, the black box 240 that is
28 created will have one or more cryptographic keys 248 associated therewith and hidden

1 therein. Once black box 240 has been created, black box server 304 transmits black box
2 240 via wide-area network 52 to remote computer 49 for installation thereon.

3 It is particular advantageous to transmit hardware ID 224 and black box
4 240 via a network such as wide-area network 52 or (if applicable) local area network
5 51, because this allows a black box 240 to be obtained in "real-time" (e.g., the entire
6 transaction of requesting, creating, and receiving a black box may take only seconds).
7 However, it will be appreciated that the use of a network is merely exemplary and not
8 limited of the means by which information can be exchanged between remote computer
9 49 and black box server 304. For example, if hardware ID 224 is deemed to be
10 sensitive information such that there is no network 51 or 52 over which it can be
11 transmitted with sufficient privacy, then remote computer 49 may store its hardware
12 ID(s) 224 on magnetic disk 29 or optical disk 31 for physical delivery to black box
13 server 304, or hardware ID 224 may be printed on paper and then physically delivered
14 to the operator of black box server 304 for entry using keyboard 40. Moreover, the
15 black box 240 created by black box server 304 may be delivered on magnetic disk 29 or
16 optical disk 31 for installation on remote computer 49; or it may be delivered via wide-
17 area network 51 even if hardware ID 224 was not delivered to black box server 304 in
18 that way.

19 Black Box Generator 37a

20 FIG. 4 shows the detail of black box generator 37a. Black box generator
21 37a comprises a random number generator 404, a code database 408, a key generator /
22 key database 448, a code generator 440, a compiler 436, and a postprocessor 452.
23 Random number generator 404 generates random (or pseudo-random) numbers 432,
24 which are used in the black box generation process as described below. Systems and
25 methods for generating pseudo-random numbers are known in the art and are therefore
26 not described herein. Code database 408 contains "diversionary code" for use by
27 diversionary code generator 416. Code database 408 may also contain templates for use
28 by cryptographic code generator 412. The nature and use of the code contained in code

1 database 408 is further described below. Key generator / key database 448 is an object
2 that either generates cryptographic keys 248 to be installed in black box 240, or that
3 stores previously-generated cryptographic keys 248 for installation in black box 240.
4 Methods of generating cryptographic keys are known in the art, and thus are not
5 provided herein. Moreover, methods of implementing a database such as code database
6 408 or key database 448 are also known in the art and thus are not provided herein.
7 Code generator, as more particularly discussed below, generates the code that
8 implements the individualized black box 240. Code generator 440 may produce code in
9 a source language (such as C or C++), in which case compiler 436 is used to compile
10 the code. If code generator 440 generates executable code directly, then compiler 436
11 may be omitted. Postprocessor 452, as more particularly discussed below, perfects
12 certain code-obfuscation techniques that are specified at the source level by code
13 generator 440 but cannot actually be performed and/or implemented until executable
14 code has been produced.

15 Code generator 440 generates the code for black box 240. Recalling that
16 black box 240 is "individualized" for remote computer 49, code generator 440 accepts
17 as input the hardware ID 224, which is received by black box server 304 at the time
18 that remote computer 49 requests an individualized black box 240. Code generator 440
19 generates the code for the black box based on hardware ID 224 – that is, the code that
20 code generator 440 includes in black box 240 is at least partly determined by hardware
21 ID 224, such that, all other things being equal, different hardware IDs 224 cause code
22 generator 440 to produce different code. Code generator 440 may also accept as input
23 one or more random numbers 432, and the code produced by code generator 440 may
24 be based on random number(s) 432 in the sense that different random number(s) cause
25 code generator 440 to produce different code. Preferably, code generator 440 accepts as
26 input both hardware ID(s) 224 and random number(s) 432, and produces code based
27 both on such hardware ID(s) 224 and such random number(s) 432. However, it will be
28 appreciated that code generator 440 could use one but not the other, or could use some

1 other value as input. It is particularly advantageous, however, for code generator 440 to
2 use both hardware ID(s) 224 and random number(s) 432, because this has the effect of:
3 (a) producing a black box 240 with code that is at least partly "random"; and (b)
4 allowing code to be included in black box 240 that binds black box 240 to hardware
5 ID(s) 224. The "randomness" aspect of the code in black box 240 is advantageous
6 because it helps to ensure that if a first black box 240 is successfully "hacked," the
7 techniques used to "hack" the first black box 240 are not easily reused on a second
8 black box 240 that has been generated with a different random number 432. The aspect
9 of the code for black box 240 being bound to hardware ID(s) 224 is advantageous
10 because it tends to make black box 240 resistant to portability.

11 Code generator 440 contains various components that handle different
12 aspects of the code generation process. The components preferably include: a
13 cryptographic code generator 412, a diversionary code generator 416, a healing code
14 generator 420, an obfuscating code generator 424, and a code reorganizer 428. In FIG.
15 4, these components are depicted as separate components of code generator 440.
16 However, it will be appreciated by those skilled in the art that the implementation
17 depicted is merely exemplary, as a single component could perform one or more of the
18 various functions described. Moreover, code generator 440 could contain a subset of
19 the components depicted, or additional components. Additionally, while separate code
20 generation and/or code manipulation components are depicted, it should be appreciated
21 that the code generated by these components need not be or remain separate; the
22 generated code can be interleaved or combined in any manner. For example,
23 cryptographic code generator 412 may create a first set of code, and diversionary code
24 generator 416 may create a second set of code, where the first and second sets do not
25 necessarily remain as separate contiguous blocks but may be woven together in any
26 manner. These variations and other may be effected without departing from the spirit
27 and scope of the invention.

28 The exemplary elements depicted in FIG. 4 are each discussed below.

Cryptographic Code Generator 412

Cryptographic code generator 412 generates the code for inclusion in black box 240 that applies cryptographic keys 248. It will be recalled that a function of the exemplary black box 240 generated by code generator 440 is to use cryptographic key(s) 248 to perform decryption and authentication services while hiding key(s) 248 from the operator of remote computer 49 on which black box 240 is to be installed. Code generator 440 obtains these cryptographic keys from key generator / key database 448 and creates code that hides them in black box. The keys 248 obtained from key generator / key database 448 may be any type or size of cryptographic keys for use with any type of cryptographic algorithm. Preferably, cryptographic keys 248 include asymmetric or "public/private" key pairs for use with public/private key encryption/decryption and authentication algorithms. One non-limiting and preferred example of a public/private key algorithm for encryption/decryption and/or authentication is the RSA algorithm described in U.S. Patent No. 4,405,829 (Rivest, et al.), although it will be appreciated by those skilled in the art that other public/private key algorithms could be used. Keys 248 are "hidden" in black box 240 in the sense that they are never actually represented in numerical form in black box 240. Instead, black box 240 contains code that performs actions that are functionally equivalent to those that would be performed by the relevant cryptographic algorithm if keys 248 were available. Regarding key "hiding" techniques, see generally *A. Shamir & N. van Someren, "Playing Hide and Seek with Keys."*

An understanding of how cryptographic code generator 412 can generate code that applies keys 248 where the code does not have access to keys 248 begins with the mathematical principle that some computations involving a given number can be performed without directly using the number, but merely by using certain properties of the number. For example, one can determine whether a decimal integer is even or odd without knowing its entire value, but merely by knowing its least significant digit: a decimal integer is even if and only if its least significant digit is 0, 2, 4, 6, or 8 (or, in

1 the case of a binary integer, if and only if its least significant bit is 0). One can
2 determine whether a number is negative or non-negative without examining the entire
3 number, but merely by examining the sign of the number. The number's least
4 significant digit (or bit) and its sign are "properties" of a number, which, in the
5 foregoing examples, may be all that one needs to know about the number in order to
6 compute the information desired. Thus, a program might receive a secret number as
7 input, where the program performs a first action if the number is negative and a second
8 action if the number is non-negative. In this example, the program could be constructed
9 to store only the sign bit of the number. Alternatively, instead of storing any
10 information about the number, the program could read the sign bit and then represent it
11 in memory by dynamically creating code that (non-obviously) produces either 0 or 1
12 depending on what the sign bit is, where any portion of the program that needs to know
13 the sign bit simply executes the dynamically-created code instead of retrieving the sign
14 bit from memory. This is a simple example of how a program can be constructed to use
15 a number without storing the number in memory or otherwise exposing the number to
16 discovery by a user.

17 Analogously, cryptographic code generator 412 of the present invention
18 makes use of mathematical properties of a particular cryptographic key 248, and creates
19 code that computes the decrypted message that results from applying key 248 to a given
20 ciphertext message without actually using key 248 itself or otherwise requiring access
21 to key 248. Cryptographic code generator 412 similarly creates code that uses key 248
22 to validate cryptographic signatures – again, without requiring access to key 248. An
23 explanation of the mathematics used in generating code to apply key 248 without
24 accessing a copy of key 248 is provided in Appendix A below. In addition to the
25 technique discussed in Appendix A, an additional key-hiding technique that may be
26 used is to embed a random number in the code for black box 240, and to use that
27 random number together with hardware ID 224 as a seed for a pseudo-random number
28 generator; as the pseudo-random number generator produces numbers, bytes of a

1 particular key may be plucked out of the number stream as needed by the cryptographic
2 calculation.

3 FIG. 5 shows the detail of an exemplary cryptographic code generator
4 412. Cryptographic code generator includes a key analysis module 504 and a code
5 producing module 508 that produces the actual code that applies key 248. The key 248
6 obtained from key generator / key database 448 is received by cryptographic code
7 generator 412 for analysis by key analysis module 504. Key analysis module 504
8 performs an analysis of key 248 in light of both its numerical properties and the
9 cryptographic algorithm that will be used to apply it to a ciphertext message (or, in the
10 case of authentication, to apply it to a digital signature). Key analysis module may
11 identifies one or more properties or "attributes" of key 248 and produces or identifies
12 one or more actions and/or functions 512 that would be performed by a cryptographic
13 algorithm in the course of applying a key 248 having the identified attributes. For
14 example, suppose that key analysis module 504 determines that the forty-second bit in
15 the binary representation of key 248 is a one (i.e., "on"), and key analysis module 504
16 is programmed with information that a particular action/function 512 is always
17 performed by a particular cryptographic algorithm whenever that algorithm applies a
18 key whose forty-second bit is a one. This action/function 512 can be identified by key
19 analysis module 504 and provided as input to code producing module 508. Key analysis
20 module 508 may identify any number of attributes of key 248, and may provide all of
21 the actions/functions 512 corresponding to these attributes to code producing module
22 508.

23 Code producing module 508 receives the actions/functions 512 that it
24 will need to perform in order to apply key 248. It will be appreciated by those skilled in
25 the art that it is possible to write a large – possibly infinite – variety of code that
26 performs a given action or function 512. It will moreover be appreciated by those
27 skilled in the art that, given a particular action or function 512, it is possible to generate
28 different code to perform that function where the particular code generated is based on

1 factors other than the action or function itself. In the example shown in FIG. 5, these
2 "other factors" include hardware ID 224 and/or random number 432. Specifically, for
3 each action or function 512, code producing module 508 produces code that performs
4 such action or function 512, where the particular code produced is based on hardware
5 ID 224 and/or random number 432.

6 Code producing module 508 can produced code based on hardware ID
7 224 and/or random number 432 in two ways. First, code database 408 may contain a
8 plurality of alternative programs that perform the same action or function 512. For
9 example, code database 408 may contain five different programs that perform a certain
10 aspect of exponentiation, with each of the programs having been written by a human
11 being and stored in code database 408. In this case, code producing module 508 may
12 select a particular one of these five different programs based on the value of hardware
13 ID 224 and/or random number 432. Second, code producing module 508 may generate
14 code to perform an action or function 512 without the need to resort to a selection of
15 human written code, but rather may write the code itself on an as-needed basis.
16 Techniques for automated generation of code by machine are known in the art and
17 therefore are not provided herein. Code producing module 508 may use automatic code
18 generation techniques to generate code, where the particular code produced is based on
19 hardware ID 224 and/or random number 432.

20 In addition to using hardware ID 224 as a parameter that merely
21 determines the particular code used to perform certain actions or functions 512, code
22 producing module may also produce code that specifically binds the produced code to a
23 remote computer 49 having hardware ID 224. For example, certain critical values
24 stored in the code could be increased by the value of hardware ID 224, and code could
25 be created to retrieve hardware ID 224 directly from the hardware of remote computer
26 49 and subtract the retrieved hardware ID 224 from those stored values, where the
27 critical value itself is the result of the subtraction. In this event, the code produced will
28 not work (or will behave unexpectedly) unless the machine on which it is running

1 provides access to the correct hardware ID 224. There are a multitude of possibilities as
2 to how code may be built to rely on hardware ID 224 or otherwise to bind the code to
3 hardware ID 224, and any of those possibilities may be used without departing from the
4 spirit and scope of the invention.

5 Diversiónary Code Generator 416

6 Returning now to FIG. 4, code generator 440 includes a diversionary
7 code generator 416. Diversiónary code generator 416 adds to black box 240
8 "diversionary code." The code produced by diversionary code generator 416 is
9 "diversionary" in the sense that it performs computations – possibly an enormous
10 quantity of computations – which preferably achieve no result of any importance to
11 black box 240's function(s) of applying cryptographic keys 248 to encrypted
12 information 204 or authenticatable data 212. While execution of the code produced by
13 diversionary code generator 416 may produce no meaningful result, the code itself
14 serves a purpose: it serves to confound attempts by hackers to analyze the code of black
15 box 240. If one (e.g., a "hacker") begins with no knowledge of how black box 240 will
16 perform its function, it is difficult or impossible to look at a particular piece of the code
17 of black box 240 and determine whether it is the part of the code that performs sensitive
18 functions, or merely the red herring produced by diversionary code generator 412. The
19 code produced by diversionary code generator 412 may appear to be producing
20 intermediate results used in the course of performing sensitive functions, where those
21 results may never actually be used in the cryptographic process. Thus, one who is
22 attempting to analyze the code of black box 240 will have to wade through (possibly)
23 enormous amounts of code before finding the crucial part of the code that performs
24 sensitive functions.

25 The code added to black box 240 by diversionary code generator 416
26 may be stored in code database 408. For example, code database may contain a
27 gigabyte of computationally-intensive code, and diversionary code generator 416 may
28 retrieve, for example, ten megabytes of this code for inclusion in black box 240. The

1 particular ten megabytes of code retrieved may be based, for example, on hardware ID
2 224 and/or random number 432. The code stored in code database 408 for use by
3 diversionary code generator 416 may, for example, be obtained from computationally-
4 intensive programs that were written for purposes otherwise unrelated to black box 240.
5 Although any code may be used, the code is preferably computationally-intensive, and
6 preferably chosen to appear to a hacker as if it is doing something of importance to the
7 cryptographic functions that black box 240 implements. For example, inasmuch as
8 public/private key cryptographic algorithms rely heavily on exponentiation, code could
9 be chosen from programs that perform mathematical computations involving large
10 amounts of exponentiation. Alternatively, instead of using code stored in code database
11 408, diversionary code generator 416 could programmatically generate the diversionary
12 code.

13 The code produced by diversionary code generator 412 may be included
14 in black box 240 in a variety of ways. For example, the code that actually implements
15 the cryptographic process could be scattered throughout the diversionary code. As
16 another variation, instead of making the code that implements the cryptographic process
17 completely non-dependent on the diversionary code, diversionary code could comprise
18 computationally intensive code that always (but non-obviously) produces the number
19 one and stores it in a register (although the code may appear to one unfamiliar with it as
20 if it might be capable of producing some other value). The register could then be
21 multiplied by a value used during the cryptographic process. This may have the
22 beneficial effects of (a) making it appear to a hacker as if the diversionary code is
23 actually performing some computation of use in the cryptographic process, thereby
24 making it difficult to distinguish the diversionary code from the "real" code, and (b)
25 causing the black box 240 to produce unexpected results if the diversionary code is
26 modified (e.g. if it is modified in such a way that it no longer produces the number
27 one).

Healing Code Generator 420

Code generator 440 includes a healing code generator 420. Healing code generator 420 creates code that uses “error-correction” principles to detect “patches” (i.e., additions or modifications to the code of black box 240 that were not part of that code as originally constituted). Healing code generator 420 may also create code that can be used to “repair” those patches by dynamically “re-modifying” the code of black box 240 to restore it to its original configuration (or to some intermediate state that exists (or should exist) during the execution of the black box 240 program).

In a conventional context, error-correction principles are employed to correct errors brought about by noise during data transmission. Application of these error correction principles to executable code is based on the notion that executable code, like any other data, is merely a sequence of bits. Therefore, minor modifications to the code (or any other data) can be detected and corrected – whether the modifications are brought about innocently by random noise, or nefariously by deliberate attacks on the code (or data). In order to create healing code, healing code generator 420 may incorporate one or more error correcting codes into the code for black box 240 as part of the individualization process. These error correcting codes (ECCs) may include linear codes (such as Hamming codes), cyclic codes (such as BCH codes) and Reed Muller codes. The code generated by healing code generator 420 can detect tampering based on parity checks and on the error syndrome that is calculated based on code and data in black box 240. Based on the calculated error syndrome, the code generated by healing code generator 420 may either correct the modification to the code, or effectively destroy the executing code by effecting other changes to the code that will cause incorrect operation of the black box 240 program. The code generated by healing code generator 420 may be designed to detect/heal tampering both as to the initial state of the black box 240 program (i.e., the state of the black box 240 program immediately after loading it into memory) and/or the running state of the program (i.e., the state of the execution space of the program at some intermediate point during its

1 execution). Additionally, as part of the individualization process, healing code
2 generator 420 may create error correcting code based on a random number and/or
3 hardware ID 224.

4 Obfuscating Code Generator 424

5 Code generator 440 includes an obfuscating code generator 424.
6 Obfuscating code generator 424 creates code that complicates examination and
7 modification of the functions performed by black box 240. It will be appreciated that
8 code produced by other components of code generator 440 is already considerably
9 resistant to analysis. As noted above, the code that applies key 248 resists discovery of
10 the key due to the fact that the key is never stored. Additionally, the "diversionary"
11 code makes it difficult to discover which code is performing the sensitive cryptographic
12 functions. Obfuscating code generator 424 adds an additional layer of protection to
13 black box 240.

14 Obfuscating code generator 424 complicates analysis and/or modification
15 of black box 240 by such techniques as: encrypting portions of the code of black box
16 240; introducing "integrity checks" into black box 240; and/or obfuscating execution of
17 code segments (e.g., by loading the code into scattered, randomly-located portions of
18 memory for execution). Obfuscating code generator 424 selects portions of the code of
19 black box 240 to be protected by encryption, integrity checks, and obfuscated
20 execution. Preferably, the encrypted portions include the portions of the code that
21 perform the most security-sensitive cryptographic functions performed by black box
22 240, but also include other portions in order to complicate the identification of the truly
23 crucial portions of the code. (I.e., if only the most sensitive portions were encrypted,
24 then it would be easy for a hacker to locate the most sensitive portions simply by
25 looking for encrypted code; therefore, these portions can be hidden by encrypting
26 unimportant portions of the code including, but not limited to, the diversionary code
27 produced by diversionary code generator 416.) Additionally, obfuscating code

1 generator 424 selects segments of the black box code to be preserved by integrity
2 checks and obfuscated execution.

3 As noted above, the code generated by code generator 440 is preferably
4 in a source language, such as C or C++. One method of introducing the above-
5 mentioned security protection into code that is created in a source language is to use a
6 software development tool such as a "secure coprocessor simulation" ("SCP") toolkit.
7 The following is a brief description of an SCP system, which is an exemplary tool used
8 by obfuscating code generator 424 to introduce protections into the code of black box
9 240.

10 SCP: Overview

11 The SCP system works at the source-code level by creating a number of
12 classes called SCPs, each of which supports an API for security operations such as
13 encryption, decryption, integrity verification, and obfuscated execution. Operations
14 such as encryption and checksum computation handle relocatable words, so that the
15 SCP system may work whether or not code is relocated (e.g., as with DLLs).

16 To protect an application with the SCP system, one must have access to
17 the application source code. Although the entity with access to the source code may be
18 a human program, this need not be the case. Specifically, an automatic code generator,
19 such as the code generator 440 used to create black box 240, may also use the SCP
20 system to insert protections into code. The entity (i.e., human programmer or code
21 generator) performs the following steps to use the SCP system to protect code:

- 22 1. SCP macros and labels are inserted into the source file(s) to be protected.

23 These macros and labels indicate code sections to be protected by such
24 measures as encryption, integrity verification, and obfuscated execution.

- 25 2. Places in the source code are selected where functions such as obfuscated
26 execution, integrity verification, decryption, and other SCP functions are to
27 be performed. SCP macros and calls are inserted into the source code at
28 these selected places to perform these actions.

1 3. SCP files are added to the application and/or compiled with the application
2 source code to yield an application executable (e.g., a .EXE or .DLL file).

3 4. A post-processing tool (e.g., postprocessor 452) is run on the application
4 executable file. The post-processing tool may locate sections of the
5 executable file to encrypt, checksum, or otherwise process in the executable.
6 The post-processing tool may also randomly (or pseudo-randomly) generate
7 cryptographic keys, encrypt the code, scatter some keys throughout the
8 executable file's code (i.e., ".text") section, and perform other setup
9 actions.

10 The follow sections describe exemplary features of the SCP system.

11 SCP: Code-integrity verification

12 A programmer (or code generator 440) may checksum any number of
13 (possibly overlapping) application segments at any time during execution. SCP provides
14 two verification methods, which are accessed via a macro-based API that generates
15 appropriate functions and data in the protected application's code (.text) section. The
16 SCP system may, for example, assume that only one code section exists and that its
17 name is ".text," as is typically the case with standard WIN32 executables. Verification
18 checks can be inlined (as described below) to avoid exposing boolean return values and
19 single points of attack.

20 Method 1: This method works with overlapping segments. The
21 programmer or code generator marks a region to be verified by inserting the macros:

22 BEGIN_VERIFIED_SEGMENT(ID1, ID2)

23 and

24 END_VERIFIED_SEGMENT(ID1, ID2)

25 inside functions, and adding the macro

26 VERIFIED_SEGMENT_REF(ID1, ID2)

27 outside of functions. The former two macros mark regions, and the latter macro creates
28 a "segment-reference function" that returns the region's addresses and checksum. The

1 variable (ID1, ID2) is a unique two-byte segment identifier, and checksums are pre-
2 computed in order of the ID1 values of regions; this is to allow cross-verification of
3 code regions by ensuring a definite order of checksum computation and storage in the
4 code section. To verify a region, the programmer or code generator inserts a function
5 such as

6 SCP.VerifySeg(VerifiedSegmentRefID1_ID2),

7 which returns a boolean value and makes available both the valid checksum and the
8 computed checksum (which must match if the code segment is intact). Any SCP object
9 can verify any segment.

10 Method 2: The second method works only with non-overlapping
11 segments. The programmer or code generator specifies a section to be verified using
12 the macros

13 BEGIN_VERIFIED_SECTION(Section, ID1, ID2)

14 and

15 END_VERIFIED_SECTION(Section, ID1, ID2)

16 which must be placed outside of functions. The variables (Section, ID1, ID2) specify a
17 unique section name and a pair of identifier bytes. The programmer or generator can
18 verify a section by inserting into the source code a function call such as

19 SCP.VerifyAppSection(BeginSectionID1_ID2, EndSectionID1_ID2)

20 both to obtain a boolean verification value and two checksums, as above.

21 It should be noted that integrity checks can be performed in a way that is
22 dependent on hardware ID 224, such that black box 240 will not operate properly if it
23 is running on the wrong machine.

24 SCP: secret scattering

25 Method 2 above uses cryptographic keys scattered using a subset-sum
26 technique. Each key corresponds to a short string used to compute indices into a
27 pseudo-random array of bytes in the code section, retrieve the bytes specified by the
28 indices, and combine these bytes into the actual key.

SCP: Obfuscated function execution

Using labels and macros, the programmer or code generator may subdivide a function into any number of blocks that the SCP system encrypts. When the function runs, SCP preferably uses multiple threads to decrypt each block into a random memory area while executing another block concurrently. Code run in this manner is preferably self-relocatable. For example, one way to implement self-relocatable code on the Intel x86 platform is to make all function calls via function pointers. Alternatively, code could be relocatable by means of dynamically adjusting call offsets during execution.

SCP: Code encryption and decryption

The programmer or code generator can specify any number of code regions to be encrypted by enclosing them within the macros

BEGIN_ENCRYPTED_SEGMENT(ID1, ID2)

and

END_ENCRYPTED_SEGMENT(ID1, ID2)

and adding the macro

ENCRYPTED_SEGMENT_REF(ID1, ID2)

outside encrypted regions. These macros and the previously described macros for verified segments serve similar purposes. If a verified region with the identifier (ID1, ID2) exists, its checksum is used as the encryption key for the corresponding encrypted segment (i.e., the one with identifier (ID1, ID2)). The programmer or code generator calls for a segment to be decrypted prior to its execution (by inserting appropriate SCP call(s)), and can then re-encrypt the segment (also using appropriate SCP call(s)). Encrypted segments may overlap, and may preferably be encrypted based on their order of appearance in the source code.

SCP: Probabilistic checking

Each SCP class has its own pseudo-random-number generator that can be used to perform security actions such as integrity verification only with certain

1 probabilities. Additionally, SCP macros may be available that produce pseudo-random
2 generators inline for this function, as well as for any other function that requires
3 pseudo-random numbers.

4 SCP: Boolean-check obfuscation

5 Several SCP macros provide means of hiding boolean verification of
6 checksums. In particular, SCP macros or objects may use checksums to mangle stack
7 and data, and compute jump addresses and indices into simple jump tables, so that
8 boolean checks become implicit and non-obvious.

9 SCP: Inlining

10 To avoid single points of attack, SCP provides macros for inline integrity
11 checks and pseudo-random generators. These macros may essentially duplicate the code
12 from the SCP segment-verification functions.

13 Code Reorganizer 428

14 Code generator 440 may include a code reorganizer 428. Code
15 reorganizer 428 reorders or resequences the code that has been produced by the other
16 components of code generator 440. For example, code reorganizer 428 may move
17 instructions in the code into a different sequential order, and then add jumps (such as
18 "conditional" jumps that test for a condition that is always, but non-obviously, true.

19 Other Features of Code Generator 440

20 Code generator 440 has some features that cut across the exemplary
21 components depicted in FIG. 4. For example, code generator 440 may create different
22 code to perform the same function within black box 240. This technique has the effect
23 of obfuscating the function being performed, since one must perform an analysis of the
24 different code segments in order to determine whether they are functionally equivalent.
25 Additionally, the differing code segments may be "inlined," so as to guard against
26 single point attacks. Other techniques that may be used by code generator 440 including
27 introducing timing loops (to detect "hijacking" of calls by an outside program), or
28 "interleaving" data, state and code by hashing and jumping on the result (which may

1 have the effect of making local modification more difficult). Code generator 440 may
2 create code that implements techniques aimed at detecting observation of black box 240
3 by a kernel-level debugger; such techniques may include changing debug registers,
4 checking to see if debug registers were changed, trapping randomly and examining
5 stack addresses. Moreover, code generator 440 may create non-sensitive code for black
6 box 240, such as routine code that performs the startup and shutdown of black box 240,
7 or that performs routine "housekeeping" tasks (although this "routine" code may still
8 be protected by techniques such as those introduced by obfuscating code generator
9 424).

10 Code generator 440 may also produce "diversionary data" for inclusion
11 in black box 240. Like the "diversionary code" discussed above, diversionary data is
12 not of direct relevance to the functions that black box 240 performs, but serves to
13 complicate analysis of black box 240. Diversionary data may include data that is never
14 used, data that is used exclusively by diversionary code, or as input to functions which
15 are designed to nullify the effect of the diversionary data. An additional feature of the
16 diversionary data is that, since its quantity can be varied, it has the effect of shifting the
17 address of black box 240's useful code between different instances of black box 240,
18 which makes it less likely that a successful attack on one black box will aid an attack on
19 another black box.

20 Compiler 436

21 Compiler 436 is used when code generator 440 creates code in a source
22 language instead of machine-executable code. For example, in an exemplary
23 embodiment code generator 440 creates code in the C++ programming language, and
24 compiler 436 is or comprises a C++ compiler. Compiler 436 receives the source code
25 produced by code generator 440 and converts it into executable code. Compilers are
26 well-known in the art and commercially available, and therefore are not discussed in
27 detail herein. Compiler 436 may be a specially-configured compiler included in black

1 box generator 37a, or it may be a general purpose compiler that is separate from black
2 box generator 37a.

3 Postprocessor 452

4 The SCP tool discussed above provides a convenient means to specify
5 certain code security techniques, such as inline code encryption, integrity verification
6 (e.g., checksums), and obfuscated execution of code. However, certain aspects of these
7 techniques cannot be performed directly on source code. For example, it is not possible
8 to encrypt source code prior to compilation. Instead, portions of the executable code to
9 be encrypted can be delimited in the source file, but the actual encryption must await
10 creation of the executable code. Similarly, segments to be protected by integrity checks
11 can be delimited in the source, but the actual creation of the integrity value must await
12 creation of the executable, since it is not possible to obtain a hash of the executable
13 code (by means of which integrity will be verified) until the source code has been
14 compiled. Postprocessor 452 performs these functions.

15 Postprocessor 452 performs, for example, the functions of: encrypting
16 the code specified for encryption, computing integrity value, generating cryptographic
17 keys for the encryption of code, and storing integrity values and code decryption keys
18 in the executable file.

19 Process of Creating an Individualized Black Box

20 Referring now to FIG. 6, an exemplary process is shown by way of
21 which black box server 304 creates and provides a secure repository that is
22 individualized for remote computer 49, such as exemplary black box 240. First, at step
23 601, black box server 304 receives the hardware ID 224 associated with remote
24 computer 49. As previously noted, hardware ID 224 may be received via a network,
25 such as wide-area network 52, in the form of a request for a black box 240, where the
26 request comes from remote computer 49. Alternatively, hardware ID 224 may be
27 received by other means, such as physical delivery on a digital medium (e.g., magnetic
28 disk 29, optical disk 31, etc.), or on paper for entry by keyboard. Black box server 304

1 may use the process depicted in FIG. 6 to create an individualized black box regardless
2 of how it receives hardware ID 224. Moreover, it will be appreciated by those skilled
3 in the art that hardware ID 224 is merely exemplary of the type of information that
4 supports the creation of a black box 240 individualized for remote computer 49. For
5 example, instead of hardware ID 224, black box server 304 may receive the serial
6 number of an operating system running on remote computer 49. Any type of
7 information will suffice, provided that it identifies remote computer 49, or is in some
8 way related to remote computer 49 or to the environment present on remote computer
9 49 (e.g., the serial number of the operating system installed on remote computer 49).

10 At step 602, black box server creates the executable code for an
11 individualized black box 240, where the code created is at least partly based on
12 hardware ID 224. As discussed above, the process of creating this executable code
13 may, for example, be performed by black box generator 37a. An exemplary process by
14 which step 602 performed is shown in FIG. 7 and discussed in detail below.

15 At step 603, black box server 304 provides the individualized black box
16 240 to remote computer 49. Black box server 304 may provide black box 240 via, for
17 example, wide-area network 52, which may be the same network over which black box
18 server 304 received hardware ID 224. Alternatively, black box server 304 may provide
19 black box 240 to remote computer 49 by physical delivery of magnetic disk 29 or
20 optical disk 31, or by any other means by which the executable file(s) in which black
21 box 240 is contained may be communicated from one computer to another.

22 FIG. 7 shows an exemplary process for performing the code creation
23 step 602 depicted in FIG. 6. The process shown in FIG. 7 may, for example, be carried
24 out by black box generator 37a (shown in FIG. 4). At step 701, black box generator
25 37a obtains a cryptographic key 248 (or possibly several cryptographic keys) to be
26 hidden in black box 240. Cryptographic key 248 may, for example, be obtained from
27 key generator / key database 448.

1 At step 702, black box generator 37a analyzes the newly obtained
2 cryptographic key 248 in order to identify one or more actions that would be performed
3 in the course of using a given cryptographic algorithm to apply cryptographic key 248
4 to data. The cryptographic algorithm may either be a decryption algorithm (that uses
5 cryptographic key 248 to convert ciphertext into cleartext), or an authentication
6 algorithm (that uses cryptographic key 248 to verify that a particular data was created
7 by the holder of a particular secret). The particular cryptographic algorithm by which
8 cryptographic key 248 will be applied is be taken into account when identifying actions
9 will be performed in the course of applying cryptographic key 248. As one example,
10 step 702 may include the act of using key analysis module 504 of cryptographic code
11 generator 412 to identify actions/functions 512, as depicted in FIG. 5.

12 At step 703, black box generator 37a generates code to perform the
13 actions identified at step 702. The code generated at step 703 is capable of applying
14 cryptographic key 248 to data, preferably without requiring access to cryptographic key
15 248 itself. The process of generating code the code may, for example, be carried out by
16 code producing module 508 of cryptographic code generator 412, depicted in FIG. 5.
17 Code may be generated programmatically, or it may be retrieved from code database
18 408. The particular code that is generated or retrieved may be at least partly based on
19 hardware ID 224 and/or random number 432. Additionally, the code produced may be
20 designed to rely in some way upon correct dynamic retrieval of hardware ID 224 from
21 the environment in which black box 240 is intended to run.

22 At step 704, black box generator 37a generates "diversionary" code for
23 inclusion in black box 240. Diversionary code may, for example, be stored in code
24 database 408. A portion of the code stored in code database 408 may then be retrieved
25 by diversionary code generator 416. The particular code retrieved for inclusion in black
26 box 240 may be based, for example, on random number 432 and/or hardware ID 224.

27 At step 705, black box generator 37a generates "healing" code, which is
28 designed to replace "patches" (i.e., unauthorized modifications to the code of black box

1 240) with the originally intended code. The healing code may be generated, for
2 example, by healing code generator 420. The particular code generated at step 705 is
3 generally based on other parts of the code to be included in black box 240, since the
4 code generated at step 705 is specifically designed to replace portions of black box 240
5 if they become modified.

6 At step 706, black box generator 37a introduces obfuscation and integrity
7 measures into the code of black box 240, such as encrypted code, integrity checks, and
8 obfuscated execution of code. Black box generator 37a may, for example, perform this
9 function by using obfuscating code generator 424 to introduce the macros and function
10 calls of the SCP system described above.

11 At step 707, black box generator 37a reorganizes the code for black box
12 240, for example by reordering or resequencing segments of the code and introducing
13 jumps to cause the code to be executed in the correct sequence. The reorganization
14 may, for example, be performed by code reorganizer 428.

15 At step 708, the code generated by black box generator 37a at steps 703
16 through 707 is optionally compiled. Preferably, the code generated at steps 703 through
17 707 would be generated in a source language such as C or C++, in which case it
18 needs to be compiled. The compilation may, for example, be performed by compiler
19 436, which is a compiler appropriate for the language in which the code was generated,
20 such as a C compiler or a C++ compiler. In an alternative embodiment in which black
21 box generator 37a generates code directly in a machine executable format (e.g., a .EXE
22 or .DLL file), then it is unnecessary to perform step 708.

23 At step 709, black box generator 37a performs postprocessing of the
24 executable code for black box 240. Step 709 may, for example, be performed by
25 postprocessor 452. As previously discussed in connection with FIG. 4, there are some
26 aspects of the code obfuscation and integrity measures introduced at step 706 that
27 cannot be perfected at the source level, but must wait until the executable code has been
28 produced. For example, one obfuscation measure is to encrypt executable code inline,

1 and decrypt it prior to its use. It is not possible to encrypt the executable code until it
2 has been produced (and it is not possible to encrypt the source code because it will not
3 compile). Moreover, it is not possible to create the integrity values (e.g., hashes,
4 checksums, etc.) that are used to implement integrity checks because these measures,
5 too, require encrypted code. Thus, at step 709 postprocessing of the executable code
6 produced at step 708 is performed. Specifically, sections of code delimited for
7 encryption are encrypted, and sections of code marked for integrity checks are hashed.
8 The keys used for encrypted code may be selected at step 709. The keys may be based
9 in part on hardware ID 224 and/or random number 432. The particular method of
10 performing an integrity check may be based in part on hardware ID 432.

11 Once postprocessing is complete, black box generator 37a proceeds to
12 step 603 shown in FIG. 6. It will be appreciated that, while the steps of FIG. 7 are
13 depicted as taking place in a certain order, certain of the steps shown may take place in
14 a different order. For example, the code generated at step 703 through 705 could be
15 generated in sequences other than that depicted in FIG. 7. The reorganization of code at
16 step 707 could take place either before or after the compilation performed at step 708.
17 Other modification may be made to the order of steps without departing from the spirit
18 and scope of the invention.

19 Example Architecture Incorporating Black Box 240 and Decoupling Interface 220

20 As noted above in connection with FIG. 2, black box 240 and application
21 program 244 may communicate either directly or through a decoupling interface 220.
22 FIG. 8 shows an exemplary architecture of a system in which a decoupling interface is
23 employed for use with black box 240 and application program 244.

24 Decoupling interface 220 addresses the issue of how black box 240 can
25 be replaced with another black box 240a, while still permitting communication and
26 authentication to take place between application program 244 on the one hand, and
27 black box 240 or 240a on the other hand. Replacement of black box 240 may become
28 necessary if black box 240 has become damaged (e.g., if a hacker has made a

1 deliberate and irreparable attempt to modify black box 240), if black box 240 has
2 become obsolete due to the development of new secure repository technology (which
3 may, for example, include either new software techniques or a hardware-based
4 repository), or if application program 244 may be usable on different types of platforms
5 having different types of secure repositories. (E.g., an open platform such as a PC
6 running one of the MICROSOFT WINDOWS 95, 98, NT, or 2000 operating systems
7 may require a different type/level of security from a closed platform such as a dedicated
8 text viewing device.) However, it will be noted that application 244 preferably should
9 be able to interface with any black box – either original or replacement – in at least two
10 ways. First, application program 244 needs to authenticate itself to the black box, since
11 the black box should not perform sensitive function for an entity who has not
12 established trustworthiness. Second, application program 244 needs to be able to
13 communicate with black box 244 in order to request performance of the sensitive
14 functions that black box 244 is designed to perform. Decoupling interface 220 provides
15 a single authentication and communication protocol that application program 244 can
16 use regardless of the black box being used, thus making the particular secure repository
17 transparent to the developer of application program 244.

18 Application program 244 and black box 240 in the exemplary
19 architecture shown communicate through decoupling interface 220. Decoupling
20 interface 220 may, for example, be or comprise an application programmer interface
21 (API) having an “initialization” call and a “bind to license” call. (Licenses are
22 explained in further detail below; briefly, a license permits the use of content protected
23 by black box 240.) The API may be provided in the form of a dynamic-link library
24 (.DLL) file that is loaded with application program 244 and executes in the process
25 (i.e., in the address space) of application program 244. Both of the calls provided by
26 the API may be exposed to application program 244 and implemented by decoupling
27 interface 220 (i.e., the calls have addresses located in the address space of application
28 program 244). Additionally, these two calls may preferably provide all of the

1 interaction that is necessary in order for application program 244 to make use of black
2 box 240. For example, the "initialization" call may perform the functions of (a)
3 authenticating application program 244 to black box 240, and (b) causing decoupling
4 interface to execute instructions that initialize black box 240 for use and allow black
5 box 240 to authenticate decoupling interface 220. A purpose of black box 240's
6 authenticating decoupling interface 220 is to establish a chain of trust. Since decoupling
7 interface 220 presents application program 244's proffer of trustworthiness to black box
8 240, black box 240 must trust decoupling interface 220 to perform this authentication
9 properly (e.g., black box 240 may need to ensure that decoupling interface 220 is not a
10 rogue DLL presenting a stolen certificate in an unauthorized context). The bind-to-
11 license call may request that black box 240 perform its sensitive function(s) for
12 application program 244. In the example depicted in FIG. 8, the sensitive function that
13 black box 240 performs is to enable the use of encrypted licensed data object 800 when
14 permitted by license 816 (as more particularly described below), and the bind-to-license
15 call may provide "enabling bits" 824 that allow application program 244 to use licensed
16 data object 800.

17 For example, application program 244 may issue an initialization call
18 (e.g., *DecouplingIF::Init()*) which is implemented by decoupling interface 220. This
19 call may cause decoupling interface 220 to execute code that starts black box 240, gives
20 it an opportunity to authenticate decoupling interface 220, and prepares a secure
21 environment that discourages modification or observation of black box 240 while black
22 box 240 is executing. It should be noted that the act of authenticating decoupling
23 interface 220 is optional, as there may be environments in which the authenticity of
24 decoupling interface 220 can be presumed from circumstance (such as purpose-built
25 devices for which all software is in the form of "firmware" installed by the
26 manufacturer). Black box 240 may take steps to prepare a secure environment, such as
27 attaching to application program 244 using the *DebugActiveProcess* system call of the
28 MICROSOFT WINDOWS 95/98/NT/2000 operating systems, in order to prevent other

1 debuggers from attaching. Preferably, decoupling interface 220 does not process the
2 bind-to-license call until decoupling interface 220 has completed the initialization
3 process with black box 240. After the initialization process is complete, application
4 program 244 may issue a bind-to-license call (e.g., *DecouplingIF::BindToLicense()*),
5 which causes decoupling interface 220 to execute code that requests black box 240 to
6 perform its sensitive functions for application program 244. In the example of FIG. 8,
7 the bind-to-license call requests that black box provide enabling bits 824. Black box 240
8 then provides enabling bits 824 to decoupling interface 220 (assuming that circumstance
9 permit), and decoupling interface 220 provides enabling bits 824 to application program
10 244. It will be appreciated that this arrangement allows application program 244 to
11 authenticate itself and request services from black box 240 without knowing the details,
12 mechanics, or communications protocols needed to interface directly with black box
13 240. In effect, decoupling interface provides a common language through which
14 application program 244 and black box 240 may communicate.

15 Still referring to FIG. 8, in the exemplary architecture shown, data
16 object 800 is a file that includes encrypted information 204, a sealed key 820, and a
17 license 816. Sealed key 820 may be the key that decrypts encrypted information 204. In
18 a preferred embodiment, sealed key 820 is a symmetric key that was used to encrypt
19 encrypted information 204. License 816 governs the use of encrypted information 204.
20 For example, license 816 may specify the rights to decrypt and/or render encrypted
21 information 204. Encrypted information 204 may, for example, be the text of a book.
22 In other examples, encrypted information 204 may be audio or video content, such as a
23 digital audio file or a digital video file. Application program 244 is a software
24 application appropriate for the nature of the information contained in data object 800.
25 For example, if encrypted information 204 is the text of a book, then application
26 program 244 may be a text-rendering application or "reader" program. If encrypted
27 information 204 is a digital audio or digital video, then application program 244 may be
28 an audio-rendering or video-rendering application.

1 The exemplary architecture of FIG. 8 includes black box 240. Black box
2 240 includes a public key 248a, and a corresponding private key 248b. As discussed
3 above, private key 248b is preferably "hidden" in black box 240, such that no copy of
4 private key 248b actually appears in black box 240, although black box 240 contains
5 the code necessary to apply private key 248b.

6 In the example of FIG. 8, there is shown a certificate 812, which has
7 associated therewith an asymmetric key pair including a public key 804 and a private
8 key 808. Private key 808 is not represented directly in certificate 812, but rather is
9 encrypted with the public key of black box 240. Sealed key 820 is "key"-sealed in data
10 object 800 with public key 804 of certificate 812 such that it can only be "unsealed"
11 with private key 808. It will be appreciated that the series of encrypted and sealed keys
12 establishes a chain of control from black box 240, to certificate 812, to sealed key 820,
13 to encrypted information 204, such that it is not possible to decrypt information 204
14 without all of these objects. Specifically, information 204 can only be accessed with
15 sealed key 820. Sealed key 820, in turn, can only be accessed with private key 808.
16 Private key 808, in turn, can only be accessed with private key 248a of black box 240.
17 This is a particularly advantageous arrangement, since it allows access to data object
18 800 to be tied to the private key 808 of certificate 812, rather than to a particular black
19 box 240. For example, if black box 240 were replaced with black box 240a (which has
20 public/private keys 248c and 248d, which are different from public/private keys 248a
21 and 248b of black box 240), data object 800 could be used by black box 240a without
22 any modification to data object 800, merely by issuing a new certificate 812a, which
23 has the same public/private keys 804 and 808, but with private key 808 being encrypted
24 with public key 248c of new black box 240a, instead of being encrypted with public key
25 248a of old black box 240.

26 Exemplary black box 240 depicted in FIG. 8 performs the function of
27 enabling application program 244 to render encrypted content 204 by providing
28 "enabling bits" 824 to application program 244. Enabling bits 824 may comprise either

1 decrypted content 208 (in which case black box 240 performs the function of applying
2 key 820 to convert encrypted content 204 into decrypted content 208), or key 820 itself
3 (in which case application program 244 uses key 820 to perform the actual decryption
4 of data object 800). Exemplary black box 240 also perform the function of validating
5 license 816 to determine whether license 816 permits use of encrypted information 240.
6 If license 816 does not permit the use of encrypted information 804 (e.g., if license 816
7 is expired, or if data object 800 is not present in an authorized environment), then black
8 box 240 does not provide enabling bits 824 to application program 244.

9 Black box 240 preferably provides enabling bits to application program
10 244 by way of decoupling interface 220. However, application program 244 and black
11 box 240 may communicate directly in accordance with aspects of the invention.
12 Decoupling interface is particularly useful in the example of FIG. 8, because it allows
13 for simple replacement of black box 240 with black box 240a without requiring any
14 modification to rendering application 244. Thus, in accordance with aspects of the
15 invention, black box 240 could be replaced with black box 240a (e.g., if black box 240
16 has become corrupted or obsolete) without any change to either rendering application
17 244 or data object 800 – merely by providing black box 240a with certificate 812a and,
18 if necessary, replacing the code of decoupling interface 220 to allow it to
19 communicate/authenticate with black box 240a. (E.g., If decoupling interface 220 is a
20 dynamically linkable library of executable code (i.e., a DLL) which links to application
21 program 244 at runtime, the DLL can be replaced with a new DLL. This replacement
22 may transparent to application program 244, since the calls made by application
23 program 244 (i.e., *Init()* and *BindToLicense()*) would simply reference the new code in
24 the new DLL.) Thus, the use of decoupling interface 220 is particularly advantageous
25 because it supports the replaceability of black box 240, and thus supports a
26 “renewable” model of security. For example, if a hardware-based repository should
27 become available, decoupling interface 220 may provide communication between the
28 not-yet-created hardware based repository and application program 244 in a manner

1 that is transparent to the developer of application program 244. As another example,
2 technology for the creation of software-based repositories could progress, thereby
3 allowing a first software-based repository to be replaced with a second software-based
4 repository having features that were not present when the first repository was created.
5 As a further exemplary feature, application program 244 may specify a "type" of
6 repository that it is able and/or willing to work with, where new repositories of that
7 "type" may not be in existence at the time that application program 244 is created, but
8 which may be developed later, thereby further supporting the "renewable" model of
9 security.

10 It is possible for black box 240 and application program 244 to
11 communicate without decoupling interface 220. For example, application program 244
12 could authenticate itself directly to black box 240 and could receive directly from black
13 box 240 the enabling bits 824 needed to use object 800. However, the use of
14 decoupling interface 220 is particularly advantageous because it supports the
15 "renewability" of black boxes, as described above. More particularly, since application
16 program 244 only needs to be able to communicate through decoupling interface 220, it
17 does not need to know any of the details about black box 240, such as how to
18 authenticate itself to black box 240, or what communication protocol is used to
19 communicate with black box 240. Thus, black box 240 could be replaced with a
20 different black box 240a, with the change being transparent to the developer of
21 application program 244. For example, black box 240a may be a different type of black
22 box for use on a different type of remote computer 49 (such as a dedicate rendering
23 device), or, as noted above, it may be a black box 240a that incorporates new security
24 technology that had not been developed at the time that black box 240 was provided to
25 remote computer 49, or it may be a hardware-based secure repository. It will be
26 appreciated that the use of decoupling interface 220 enables the use of black box 240a
27 with application program 244 and data object 800 even if black box 240a has not been

1 developed or is otherwise not yet in existence at the time that application program 244
2 and data object 800 are installed on remote computer 49.

3 Referring now to FIG. 9, there is shown an exemplary process by which
4 application program 244 uses black box 240 through decoupling interface 220. At step
5 901, application program 244 starts execution. Application program 244 may be a
6 program executing on an open platform or general purpose computer, such as remote
7 computer 49. Alternatively, application program 244 may be a program for a closed
8 platform such as purpose-built hardware (not shown). For example, application
9 program 244 may be a program that displays electronically-distributed text to a (human)
10 user. One version of application program 244 may be designed to run on general-
11 purpose open-platform remote computer 49. Another version of application program
12 244 may be designed to run on a dedicated hand-held reading device that runs no
13 software other than that necessary to display electronic text. Preferably, application
14 program 244 loads decoupling interface 220 (e.g., by linking with a DLL that contains
15 instructions which implement decoupling interface 220), in order to facilitate
16 communication with black box 240.

17 At step 902, application program issues an instruction to initialize black
18 box 240. Preferably, the initialization instruction is implemented by decoupling
19 interface 220 and takes the form of a method that decoupling interface 220 exposes to
20 application program 244. For example, the code of application program 244 may
21 include an instruction of the form (*DecouplingIF::Init()*), which executes code to
22 perform the initialization function (where the code is in the decoupling interface DLL
23 and linked to application program 244 at runtime). The call to *Init()* may, optionally, be
24 parameterized by data representing a type of black box 240 (e.g., if the developer
25 and/or provider of application program 244 has deemed that application program 244
26 should work only with a particular type of black box, decoupling interface 220 may
27 provided support for such a condition if it is specified as a parameter). Steps 903 and

1 904, described below, may be performed as part of, or in response to, the initialization
2 instruction issued at step 902.

3 At step 903, decoupling interface 220 proffers to black box 240 proof of
4 the authenticity and/or trustworthiness of application program 244. The process of
5 authenticating application program 244 generally includes obtaining a certificate from
6 application program 244 (where the certificate is signed with the private key of a
7 trusted certifying authority) and then validating the signature. The certificate may be
8 located in the code of application program 244 in a place known to the DLL that
9 implements decoupling interface 220. Black box 240 preferably incorporates the public
10 key of the certifying authority. The process of authenticating application 244 may be
11 particularly simple when the platform on which application 244 is running is a purpose
12 built device or features hardware-enforced isolation, since it may then be presumed that
13 only trusted applications 244 would be installed on such a device.

14 At step 904, black box 240 is started and, optionally, is given the
15 opportunity to prepare a secure environment, including authenticating decoupling
16 interface 220 and application program 244. The details of this process depend on the
17 particular environment in which black box 240 and application program 244 are
18 running. For example, in the case of a closed, purpose built device, the only actions
19 that may need to take place are retrieving a private key from a ROM and validating its
20 certificate.

21 In the case of "authenticated boot" environments, the process may
22 include using an operating system call to check a software image (e.g., of the
23 decoupling interface 220 DLL), checking the certificate of application program 244,
24 and isolating the code and data space of application program 244.

25 In the case of an open platform, such as where application program 244
26 runs on a typical personal computer using one of the MICROSOFT WINDOWS
27 95/98/NT/2000 operating systems, the process of authenticating decoupling interface
28 220 and preparing a secure environment may include the following actions. The

1 initialization call of decoupling interface 220 (i.e., *DecouplingIF::Init()*) may cause the
2 code in decoupling interface 220's DLL to issue an initialization call exposed by black
3 box 240 (e.g., *BlackBox::Init*). Preferably, one of the arguments to *BlackBox::Init* may
4 be a binary certificate that contains a cryptographic hash of the code of decoupling
5 interface 220 signed by the private key of a certifying authority, where the
6 corresponding public key is available to (e.g., inside) black box 240. Black Box 240
7 reads the code of decoupling interface 220, for example by using the system call
8 *ReadProcessMemory* (or a private in-process protocol that accesses the address space of
9 decoupling interface 220), and computes its hash and compares it to the one signed by
10 the certifying authority. *BlackBox::Init* may also check that the return address from its
11 *Init* call is in the execution space of decoupling interface 220. Decoupling interface 220
12 may also receive two additional certificates from black box 240: a certificate from a
13 certifying authority (which may or may not be the same certifying authority that signs
14 the certificate of decoupling interface 220) binding the name and public key of the
15 developer of application program 244, and also a binary certificate from the developer
16 of application program 244 naming the hash, security level and expiration of its code.
17 Decoupling interface 220, which is in the same address space as application program
18 244, computes the hash of the application code, verifies the certificate and compares the
19 computed hash to the signed hash. These action may be performed before decoupling
20 interface 220 returns from the initialization call issued by application program 244.

21 Additionally, during the call to *BlackBox::Init*, black box 240 preferably
22 attaches to application program 244 using the *DebugActiveProcess* system call (or
23 similar call, when an operating system other than one of the MICROSOFT WINDOWS
24 operating systems is being used), which has the advantageous effect of preventing user
25 mode debuggers from attaching to application program 244. Preferably, no sensitive
26 code within black box 240, and no key within the black box 240 (e.g., key 248) are
27 "uncovered" until this attachment is done, internal black box integrity checks are
28 performed, and the code of application program 244 and/or decoupling interface 220

1 have been authenticated. Thereafter, all sensitive data is “poked” into the application
2 process (using WriteProcessMemory, or a private in-process interface that accesses the
3 address space of the application process) and does not travel via COM, RPC or system
4 buffers.

5 Preferably, in order to establish a secure environment, system
6 components used by black box 240, decoupling interface 220, and application program
7 244 are also authenticated. Components to be authenticated may include user libraries,
8 kernel components, or any other software that has the potential to affect the security of
9 the environment in which protected actions are to be performed.

10 At the conclusion of step 904, establishment of a chain of trust between
11 black box 240 and application program 244 is complete. Black box 240 may now
12 proceed to perform sensitive functions, such as cryptographic services, for application
13 program 244.

14 At step 905, application program 244 requests services from black box
15 240. For example, black box 905 may provide or identify data object 800 and request
16 that that data object 800 be decrypted if permitted by license 816. In the case where
17 decoupling interface 220 is present, application program 244 may make this request
18 through decoupling interface 220, which then issues the necessary instructions to black
19 box 240.

20 At step 906, black box 240 performs functions necessary to validate
21 license 816 and provide application program 244 with access to data object 800. For
22 example, in the exemplary architecture depicted in FIG. 8 where license 816, key 820,
23 and content 204 are cryptographically “sealed” together, black box 240 may use its
24 (preferably hidden) private key 248b to decrypt private key 808 of certificate 812, and
25 then may use private key 808 to unseal data object 800 and read its license 816. Black
26 box 240 may then proceed to evaluate license 816, checking whatever conditions are
27 necessary to perform the evaluation. For example, license 816 may permit decryption
28 of data object 800 up to a particular date, in which case the license evaluation process

1 includes checking a (preferably trusted) calendar source. As another example, license
2 816 may permit data object 800 to be used only in a particular environment (e.g., the
3 license may permit data object 800 to be used on an isolated hardware device but not on
4 an open platform), in which case the evaluation process includes the act of identifying
5 the environment.

6 If the license permits use of data object 800, black box 240 may obtain
7 sealed key 820 and use it to provide "enabling bits" 824 to application program 244 at
8 step 907. In one exemplary embodiment, enabling bits 824 may include sealed key 820,
9 which application program 244 uses to object 800's encrypted information 204. In
10 another exemplary embodiment, black box 240 uses sealed key 820 to convert
11 encrypted information 204 into decrypted information 208 and then provides decrypted
12 information 208 to application program 244. In the latter example, the decrypted
13 information itself is the "enabling bits" 824. When decoupling interface 220 is present,
14 black box 240 preferably provides enabling bits 824 to application program 244 through
15 decoupling interface 220.

16 It is noted that the foregoing examples have been provided merely for the
17 purpose of explanation and are in no way to be construed as limiting of the present
18 invention. While the invention has been described with reference to various
19 embodiments, it is understood that the words which have been used herein are words of
20 description and illustration, rather than words of limitations. Further, although the
21 invention has been described herein with reference to particular means, materials and
22 embodiments, the invention is not intended to be limited to the particulars disclosed
23 herein; rather, the invention extends to all functionally equivalent structures, methods
24 and uses, such as are within the scope of the appended claims. Those skilled in the art,
25 having the benefit of the teachings of this specification, may effect numerous
26 modifications thereto and changes may be made without departing from the scope and
27 spirit of the invention in its aspects.

APPENDIX ABackground

Let e be the exponent to be computed.

Let e_{acc} be the exponent of the "accumulation," the number that accumulates the message raised to e .

Let e_{err} be the current error exponent.

When using the binary method to compute an addition chain, e is scanned from left to right. For each scanned bit, the accumulation is squared, which doubles its exponent (the shift stage: $e_{acc} = e_{acc} << 1$). If the scanned bit is 1, the accumulation is then multiplied by the message, which increases its exponent by 1 (the add stage: $e_{acc} = e_{acc} + 1$). This process may be characterized as shifting the MSB off e and on to e_{acc} . The add stage effectively corrects the exponent error that accumulated when the exponent was left shifted.

The m -ary method generalizes this method by shifting m bits at a time ($e_{acc} = e_{acc} << m$). At each step, the exponent error that is corrected in the add stage equals the m -bit number currently being scanned ($e_{err} = e_m$). Usually the error is resolved with one multiply using a pre-computed table of messages raised to all possible m -bit exponents.

In greater generality, when the exponent error is "resolved" by multiplying from a pre-computed table, the exponent in the table is effectively subtracted from the current exponent error (the reduction stage: $e_{err} = e_{err} - e_t$). In the binary and m -ary methods, the reduction exponent is always chosen to equal the current error, so that the reduction always yields an error of zero ($e_t = e_m$). However, a random exponent could be chosen to reduce the error as long as it was less than or equal the current error. The excess error would be carried into the next shift stage ($e_{err} = (e_{err} << m) + e_m$).

Random Addition Chains:

A library of functions (e.g., `Addchain.{h,cpp}`) can be created which describe a class that computes a random addition chain for a big number exponent. Computation of the addition chain assumes a storage mechanism, the “register store,” that holds several big numbers (typically 15) and the original message. These registers store the message raised to various 29-bit exponents as well as the number that accumulates the message raised to the total exponent (the accumulation).

The register store is primed with random exponents computed using random multiplications between already computed exponents, starting with the original message in one of the registers. While accumulating e , multiplications between random registers are performed and stored, so that “pre-computed” exponents are constantly changing. These random multiplications are tuned so that register exponents do not exceed 29 bits.

While accumulating, the exponent error is not allowed to grow more than twice as large as the largest exponent in the register store. If this limit is about to be exceeded, the error is reduced. These “required” reductions prefer larger exponents to keep the number of required multiplies small. “Optional” reductions randomly occur if there exists any register that is less than or equal the current error. The parameters for these required and optional reductions are tuned so that about 700 multiplies occur for a 512-bit exponent in addition to the 512 necessary squarings – about 3 reductions for every 2 shifts.

Since squarings are treated as multiplies, reductions are fairly unpredictable, and extraneous multiplies frequently occur, an opponent must track the contents of all the registers to determine the accumulating exponent.

Register Store:

Two categories of techniques are used to make it difficult to track the contents of a virtual register: data obfuscation and engineering tricks. Data obfuscation involves storing registers redundantly, swapping data stored in two registers, and hiding

1 data by XOR'ing it with another register. Engineering tricks help defeat common
2 techniques used for backwards engineering code, such as fall through cases on switch
3 statements or extraneous no-op calculations injected between real calculations.

4 Data Obfuscation:

5 The register store contains 66 "real" registers that are used to store data
6 for the 16 "virtual" registers used by the addition chain calculation. A virtual register
7 can keep track of up to 4 redundant copies of its data with 4 "slots" that reference real
8 registers. When data needs to be stored, a real register is randomly selected from the
9 set of unallocated registers. Real registers are released to the unallocated store when a
10 virtual register is overwritten. Dynamically allocating real registers to a virtual register
11 prevents an adversary from statically examining data stored at particular memory
12 offsets to determine the contents of the virtual register.

13 When data is needed, one of the slots is randomly selected. Duplicating
14 data for redundancy is scheduled randomly, but the process is biased to prefer data that
15 has been used recently (ex. as a source or result of a multiply). Slots are sometimes
16 released randomly, even though the data in the real register is not changed and might
17 appear to be in use. This way, a virtual register's data may actually appear in more than
18 4 real registers even though the system is only selecting data from one of the 4 slots it
19 is tracking.

20 In the process of copying data, a real register may also become XOR'd
21 with another random real register, to mask its contents. The other register may or may
22 not contain data that is being actively used. In fact, a copy operation may change the
23 XOR state of up to four other registers, in addition to the real registers being copied to
24 and from. These six registers can also swap their contents in various ways during a
25 copy via a series of XOR's. Thus, a particular copy generally involves several XOR's
26 that can change both the XOR state of several real registers and the real registers
27 referenced by virtual register slots.

Dynamic register allocation, surreptitious swapping, and XOR'ing constitutes a kind of "shell game," where data is being shuffled around and masked in memory before and after virtually every multiplication. Data stored in one register may be retrieved from an entirely different register later on. The same register may never be used to retrieve the same piece of data twice in a row. The XOR'ing makes it more difficult to track the swapping in memory.

Engineering Tricks

Static analysis of the register store's memory is deterred in three ways. Memory is allocated on the heap so techniques that examine offsets from the stack pointer do not work. Random blocks of unused space are also added between real registers so that obvious offsets cannot be used to examine a particular register. Finally, the register store is initialized with random content so that unused or unprimed registers are not obvious by examination.

The bulk of the code in rsagen.cpp is devoted to generating the `_Copy` function, and tracking the state changes that it performs as side effects. The `_Copy` function performs the copy/swap/state-change operation for the data in the register store. A single 32-bit unsigned integer is passed to `_Copy` that encodes the type of copy that is to occur. Part of this UINT encodes the 6-8 real registers involved in the copy/swap/state-change, and another part performs the actual calculation. Both of these parts are implemented with large switch statements (150+ cases) that make extensive use of fall through cases. Fall through makes it more difficult to analyze the code generated for the switch based on jumps, since 1 to 4 cases may use the same break.

The major vulnerability of the system up to this point involves a dynamic attack, where the adversary breaks out of execution at each multiply and examines the data going in and out. If the adversary knows what exponent is represented by the source data, it can compute the exponent of the result. Since the whole system starts with the original message, the adversary could backwards engineer any exponent using this attack – most notably the secret exponent.

This is deterred in two ways. First, data for the sources and result of the multiply are directly read from and written to the register store, without making an intermediate copy to stack memory. This makes it more difficult for an automatic breaker to analyze the stack frame to infer exponents. Second, a percentage of the calculations are performed using inlined copy cases and multiplies. These inlined calculations mostly occur at the beginning of the computation, but other clusters are scattered throughout the calculation randomly. With copies and multiplies mixed together, it should be more difficult to determine where exactly a multiply is occurring and therefore where to examine data to backwards engineer exponents.

Vulnerability and Possible Solution:

Currently, the greatest vulnerability of the system is still the dynamic attack. If an adversary can infer where multiplies are occurring in the inlined code, it might still backwards engineer the exponents automatically. This will be difficult with BBT code shuffling and bogus code injected into the real calculations – but it is still conceivable.

A number theory result could be used to make this more difficult. If during a copy, the code performed a surreptitious add much like the swaps and XOR's, the code could mask a source value of the multiply. After the multiply, a second number is added to the result, such that the total amount added is congruent to zero with respect to the modulus.

Given $(A+B) \bmod N = 0$, it is needed to hide the sources and result of the multiply:

$$(X*Y) \bmod N =$$

$$(X*Y + (A+B)*Y) \bmod N =$$

$$(X*Y + A*Y + B*Y) \bmod N =$$

$$((X + A) * Y + B*Y) \bmod N$$

So, add A to X before the multiply, and $(B*Y) \bmod N$ after the multiply. It is possible to add $(B*Y) \bmod N$ after the mod – as long as the 512-bit number does not overflow

1 with the add, and as long as another mod is done before the end. The copy code is
2 designed to support a surreptitious add, although this new kind of state is not
3 implemented in the state tracking code.